1. (**10 points**)
What are the possible outputs of the following program ? (Assume the new process is created successfully.)

```
#include <sys/types.h>
#include <stdio.h>
#include <unistd.h>

int main()
{
  int value = 2018;

  printf("Begin with %d\n", value);
  if(fork() <= 0) {
    value += 5;
    printf("The value is %d\n", value);
  } else {
    printf("The value is %d\n", value);
    wait(NULL);
  }
  printf("Over with %d\n", value);
}
```

2. (**30 points**)
Answer Yes or No for following questions. Each correct answer gets 5 points.
   a. Cache is useful only because it is faster than memory.
   b. With the help of virtual memory, an OS can always executes a program which size is bigger than the size of physical memory.
   c. The LRU page replacement algorithm does not always perform well.
   d. The instruction to issue a trap is not privileged.
   e. When interrupts occur, the state of the running process will transit from RUN to READY immediately.
   f. A spinlock is not useful in an uni-processor system.

3. (**10 points**)
The memory access time is 200 ns if page fault does not happen. But if page fault do happen, the memory access time is X ns. Assume the page fault rate is 0.0000025 and the effective access time does not exceed 220 ns, what is the maximum value of X?

4. (**20 points**)
There is a processor with the classical 5-stage pipeline and having 2.5 CPI with perfect cache, i.e., all cache accesses are hits. The clock frequency of the processor is 1GHz. The processor has a L1 instruction cache, and a L1 data cache. The block size is 64B for both of the caches. The cache sizes of the two caches are 32KB, and 64KB, respectively. L1 instruction cache is direct-mapped cache, and L1 data cache is a two-way set-associative cache. The memory address used by the processor is 64-bit long.
   a. (5 pts) How many tag bits does the L1 instruction and data caches respectively have?
   b. (5 pts) Assume we have a new improvement that increases the number of pipeline stages from 5 to 7 stages. The clock frequency will be increased from 1GHz to 1.2 GHz. However, the CPI with perfect cache will also be increased from 2.5 to 3. Assume our target workload does have perfect cache. Does the improvement bring any speedup? If yes, how much?
   c. (5 pts) The engineers in the company developed a new compiler that reduces 20% of instructions compared to the old one. However, the perfect cache CPI will be increased to 2.8. Should we use the new compiler?
   d. (5 pts) Another target workload has 10% of instructions that access data, and 2% of instruction cache misses, and 5% of data cache misses. Assume the cache sizes are unchanged, and L1 instruction cache is reconfigured to 4-way set-associative cache. The L1 instruction cache miss rate is reduced to 1%. However, the perfect cache CPI will be increased to 3. Does this reconfiguration bring any performance improvement?

5. (**6 points**) What are the pros and cons of having more pipeline stages?

6. (**4 points**) What kinds of workload behaviors are suitable for using Graphic Processing Units? What are your reasons?

7. (**20 points**) True or false
   a. (4 pts) Memory hierarchy is still effective even if the target workload has zero spatial and temporal locality.
   b. (4 pts) We can use instruction scheduling to hide pipeline stalls caused by pipeline hazards.
   c. (4 pts) Control hazards is caused by instructions like branch and jump.
   d. (4 pts) Since jump instructions, such as j or jal, do not have to resolve the condition like branch instructions, jump instructions do not cause control hazards.
   e. (4 pts) Fully-associative cache causes zero conflict miss.